

What is claimed is:

- 1           1. A method of dynamically modifying a cluster communication parameter in  
2 a clustered computer system, the method comprising:
  - 3           (a) initiating a cluster communication parameter modification by  
4 transmitting a message to a plurality of nodes in the clustered computer  
5 system;
  - 6           (b) locally confirming, within each node, receipt of the message by  
7 each of the plurality of nodes;
  - 8           (c) in response to confirming receipt of the message by each of the  
9 plurality of nodes, invoking a local cluster communication parameter  
10 modification operation on each node;
  - 11           (d) transmitting from each node a status of the local cluster  
12 communication parameter modification invoked on that node;
  - 13           (e) locally detecting, within each node, an unsuccessful status for the  
14 local cluster communication parameter modification on any node; and  
15           (f) in response to detecting an unsuccessful status for any node, locally  
16 undoing, in each node for which the local cluster communication operation  
17 was performed, the local cluster communication parameter modification  
18 operation performed on that node.
- 1           2. The method of claim 1, wherein the cluster communication parameter  
2 comprises a heartbeat parameter.
- 1           3. The method of claim 1, wherein the cluster communication parameter is  
2 selected from the group consisting of heartbeat message time out, heartbeat  
3 acknowledgment message time out, heartbeat frequency or interval, heartbeat failure  
4 threshold, heartbeat acknowledgment failure threshold, receive/send timer ratio,  
5 maximum fragment size, message retry timer value, maximum message retry time,  
6 send queue overflow threshold, message send window size, and combinations thereof.

1           4. The method of claim 1, wherein locally confirming receipt of the message  
2 by each of the plurality of nodes includes participating in an ACK round responsive to  
3 receipt of the message.

1           5. The method of claim 1, wherein transmitting from each node a status of the  
2 local cluster communication parameter modification invoked on that node is  
3 performed during an ACK round performed subsequent to invoking the local cluster  
4 communication parameter modification operation.

1           6. The method of claim 1, wherein transmitting the message, confirming  
2 receipt of the message, and transmitting the status are performed via multicast  
3 messages.

1           7. An apparatus, comprising:  
2               (a) a memory; and  
3               (b) a program resident in the memory, the program configured to  
4           dynamically modify a cluster communication parameter on a local node  
5           among a plurality of nodes in a clustered computer system, the program  
6           configured to locally confirm, for the local node, successful receipt of an  
7           initiation message by each of the plurality of nodes, and a status for a local  
8           cluster communication parameter modification operation performed by each of  
9           the plurality of nodes, the program further configured to undo a local cluster  
10          communication parameter modification operation performed on the local node  
11          in response to detection of an unsuccessful status for a local cluster  
12          communication parameter modification on any node.

1           8. The apparatus of claim 7, wherein the program is further configured to  
2          locally confirm receipt of an initiating message by each of the plurality of nodes.

1           9. The apparatus of claim 8, wherein the program is configured to locally  
2          confirm receipt of the initiating message by each of the plurality of nodes by  
3          participating in an ACK round responsive to receipt of the message.

1           10. The apparatus of claim 7, wherein the program is further configured to  
2          transmit from the local node a status of the local cluster communication parameter  
3          modification operation.

1           11. The apparatus of claim 10, wherein the program is configured to transmit  
2          the status during an ACK round performed subsequent to invocation of the local  
3          cluster communication parameter modification operation.

12. A clustered computer system, comprising:

(a) a plurality of nodes coupled to one another over a network; and

(b) a plurality of programs, each local to a node among the plurality of nodes, each program configured to dynamically modify a cluster communication parameter on its respective local node, each program further configured to locally confirm, for its respective local node, successful receipt of an initiation message by each of the plurality of nodes, and a status for a local cluster communication parameter modification operation performed by each of the plurality of nodes, and each program further configured to undo a local cluster communication parameter modification operation performed on its respective local node in response to detection of an unsuccessful status for a local cluster communication parameter modification on any node.

- 1           13. A program product, comprising:
- 2               (a) a program configured to dynamically modify a cluster
- 3               communication parameter on a local node among a plurality of nodes in a
- 4               clustered computer system, the program configured to locally confirm, for the
- 5               local node, successful receipt of an initiation message by each of the plurality
- 6               of nodes, and a status for a local cluster communication parameter
- 7               modification operation performed by each of the plurality of nodes, the
- 8               program further configured to undo a local cluster communication parameter
- 9               modification operation performed on the local node in response to detection of
- 10              an unsuccessful status for a local cluster communication parameter
- 11              modification on any node; and
- 12              (b) a signal bearing medium bearing the program.

- 1           14. The program product of claim 13, wherein the signal bearing medium
- 2           includes at least one of a transmission medium and a recordable medium.

- 1           15. The program product of claim 13, wherein the program is further
- 2           configured to locally confirm receipt of an initiating message by each of the plurality
- 3           of nodes.

- 1           16. The program product of claim 15, wherein the program is configured to
- 2           locally confirm receipt of the initiating message by each of the plurality of nodes by
- 3           participating in an ACK round responsive to receipt of the message.

- 1           17. The program product of claim 13, wherein the program is further
- 2           configured to transmit from the local node a status of the local cluster communication
- 3           parameter modification operation.

- 1           18. The program product of claim 17, wherein the program is configured to
- 2           transmit the status during an ACK round performed subsequent to invocation of the
- 3           local cluster communication parameter modification operation.

1           19. A method of dynamically modifying a heartbeat parameter in a node  
2 among a plurality of nodes in a clustered computer system, the plurality of nodes  
3 including first and second nodes, the first node configured to send a heartbeat  
4 message to the second node, and the second node configured to send an  
5 acknowledgment message to the first node in response to receiving the heartbeat  
6 message, the method comprising:

7                   (a) sending a heartbeat message from the first node to the second  
8 node, the heartbeat message indicating that a heartbeat parameter is to be  
9 modified; and

10                   (b) deferring modification of the heartbeat parameter in the first node  
11 until receipt of an acknowledgment message sent from the second node to the  
12 first node that indicates that the heartbeat parameter has been modified in the  
13 second node.

1           20. The method of claim 19, further comprising determining whether  
2 modifying the heartbeat parameter on the first node requires synchronization with the  
3 second node.

1           21. The method of claim 20, wherein determining whether modifying the  
2 heartbeat parameter on the first node requires synchronization with the second node  
3 further comprises determining whether the heartbeat parameter is local or global in  
4 nature.

1           22. The method of claim 19, further comprising, in response to receiving the  
2 heartbeat message with the second node, sending an acknowledgment message from  
3 the second node to the first node, the acknowledgment message indicating whether the  
4 heartbeat parameter has been modified in the second node.

1           23. The method of claim 22, wherein each of sending the heartbeat message  
2 and sending the heartbeat acknowledgment message includes accessing a heartbeat  
3 message record that includes a change request indicator, the method further  
4 comprising:

5 (a) prior to sending the heartbeat message that indicates that the  
6 heartbeat parameter is to be modified, setting the change request indicator in  
7 the heartbeat message record; and

8 (b) prior to sending the heartbeat acknowledgment message that  
9 indicates whether the heartbeat parameter has been modified in the second  
10 node, selectively setting or clearing the change request indicator in the  
11 heartbeat message record.

1 24. The method of claim 23, wherein deferring modification of the heartbeat  
2 parameter in the first node until the acknowledgment message indicates that the  
3 heartbeat parameter has been modified in the second node includes modifying the  
4 heartbeat parameter in the first node only after receiving a heartbeat acknowledgment  
5 message with a set change request indicator.

1 25. The method of claim 19, further comprising:

2 (a) modifying the heartbeat parameter in the second node; and

3 (b) modifying the heartbeat parameter in the first node after receipt of  
4 an acknowledgment message sent from the second node to the first node that  
5 indicates that the heartbeat parameter has been modified in the second node.

1           26. An apparatus, comprising:  
2                 (a) a memory; and  
3                 (b) a program resident in the memory and configured to dynamically  
4           modify a heartbeat parameter in a first node among a plurality of nodes in a  
5           clustered computer system by sending a heartbeat message to a second node  
6           among the plurality of nodes that indicates that the heartbeat parameter is to be  
7           modified and thereafter deferring modification of the heartbeat parameter in  
8           the first node only after receiving an acknowledgment message from the  
9           second node indicating that the heartbeat parameter has been modified in the  
10          second node.

1           27. The apparatus of claim 26, wherein the program is further configured to  
2          determine whether modifying the heartbeat parameter on the first node requires  
3          synchronization with the second node.

1           28. The apparatus of claim 27, wherein the program is configured to  
2          determine whether modifying the heartbeat parameter on the first node requires  
3          synchronization with the second node by determining whether the heartbeat parameter  
4          is local or global in nature.

1           29. The apparatus of claim 26, wherein the program is configured to send the  
2          heartbeat message by accessing a heartbeat message record that includes a change  
3          request indicator, and wherein the program is further configured to set the change  
4          request indicator in the heartbeat message record prior to sending the heartbeat  
5          message that indicates that the heartbeat parameter is to be modified.

1           30. The apparatus of claim 29, wherein the program is configured to defer  
2          modification of the heartbeat parameter in the first node until the acknowledgment  
3          message indicates that the heartbeat parameter has been modified in the second node  
4          by modifying the heartbeat parameter in the first node only after receiving a heartbeat  
5          acknowledgment message with a set change request indicator.



1       31. A program product, comprising:

2               (a) a program configured to dynamically modify a heartbeat parameter  
3       in a first node among a plurality of nodes in a clustered computer system by  
4       sending a heartbeat message to a second node among the plurality of nodes  
5       that indicates that the heartbeat parameter is to be modified and thereafter  
6       deferring modification of the heartbeat parameter in the first node only after  
7       receiving an acknowledgment message from the second node indicating that  
8       the heartbeat parameter has been modified in the second node; and

9               (b) a signal bearing medium bearing the program.